

Small-Bias Sets for Nonabelian Groups: Derandomizing the Alon-Roichman Theorem

Sixia Chen

Cristopher Moore

Alexander Russell

May 1, 2013

Abstract

In analogy with ε -biased sets over \mathbb{Z}_2^n , we construct explicit ε -biased sets over nonabelian finite groups G . That is, we find sets $S \subset G$ such that $\|\mathbb{E}_{x \in S} \rho(x)\| \leq \varepsilon$ for any nontrivial irreducible representation ρ . Equivalently, such sets make G 's Cayley graph an expander with eigenvalue $|\lambda| \leq \varepsilon$. The Alon-Roichman theorem shows that random sets of size $O(\log |G|/\varepsilon^2)$ suffice. For groups of the form $G = G_1 \times \cdots \times G_n$, our construction has size $\text{poly}(\max_i |G_i|, n, \varepsilon^{-1})$, and we show that a set $S \subset G^n$ considered by Meka and Zuckerman that fools read-once branching programs over G is also ε -biased in this sense. For solvable groups whose abelian quotients have constant exponent, we obtain ε -biased sets of size $(\log |G|)^{1+o(1)} \text{poly}(\varepsilon^{-1})$. Our techniques include derandomized squaring (in both the matrix product and tensor product senses) and a Chernoff-like bound on the expected norm of the product of independently random operators that may be of independent interest.

1 Introduction

Small-bias sets are useful combinatorial objects for derandomization, and are particularly well-studied over the Boolean hypercube $\{0, 1\}^n$. Specifically, if we identify the hypercube with the group \mathbb{Z}_2^n , then a *character* χ is a homomorphism from \mathbb{Z}_2^n to \mathbb{C} . We say that a set $S \subseteq \mathbb{F}_2^n$ is ε -biased if, for all characters χ ,

$$\left| \mathbb{E}_{x \in S} \chi(x) \right| \leq \varepsilon,$$

except for the trivial character $\mathbb{1}$, which is identically equal to 1. Since any character of \mathbb{F}_2^n can be written $\chi(x) = (-1)^{k \cdot x}$ where $k \in \mathbb{Z}_2^n$ is the “frequency vector,” this is equivalent to the familiar definition which demands that on any nonzero set of bits, x 's parity should be odd or even with roughly equal probability, $(1 \pm \varepsilon)/2$.

It is easy to see that ε -biased sets of size $O(n/\varepsilon^2)$ exist: random sets suffice. Moreover, several efficient deterministic constructions are known [13, 1, 3, 4] of size polynomial in n and $1/\varepsilon$. These constructions have been used to derandomize a wide variety of randomized algorithms, replacing random sampling over all of $\{0, 1\}^n$ with deterministic sampling on S (see e.g. [5]). In particular, sampling a function on an ε -biased set yields a good estimate of its expectation if its Fourier spectrum has bounded ℓ_1 norm.

The question of whether similar constructions exist for nonabelian groups has been a topic of intense interest. Given a group G , a *representation* is a homomorphism ρ from G into the group $U(d)$ of $d \times d$ unitary matrices for some $d = d_\rho$. If G is finite, then up to isomorphism there is a finite set \hat{G} of *irreducible* representations, or *irreps* for short, such that any representation σ can be written as a direct sum of irreps. These irreps form the basis for harmonic analysis over G , analogous to classic discrete Fourier analysis on abelian groups such as \mathbb{Z}_p or \mathbb{Z}_2^n .

Generalizing the standard notion from characters to matrix-valued representations, we say that a set $S \subseteq G$ is ε -biased if, for all nontrivial irreps $\rho \in \widehat{G}$,

$$\left\| \mathbb{E}_{x \in S} \rho(x) \right\| \leq \varepsilon,$$

where $\|\cdot\|$ denotes the operator norm. There is a natural connection with expander graphs. If we define a Cayley graph on G using S as a set of generators, then G becomes an expander if and only if S is ε -biased. Specifically, if M is the stochastic matrix equal to $1/|S|$ times the adjacency matrix, corresponding to the random walk where we multiply by a random element of S at each step, then M 's second eigenvalue has absolute value ε . Thus ε -biased sets S are precisely sets of generators that turn G into an expander of degree $|S|$.

The Alon-Roichman theorem [2] asserts that a uniformly random set of $O((\log |G|)/\varepsilon^2)$ group elements is ε -biased with high probability. Thus, our goal is to derandomize the Alon-Roichman theorem—finding explicit constructions of ε -biased sets of size polynomial in $\log |G|$ and $1/\varepsilon$. (For another notion of derandomizing the Alon-Roichman theorem, in time $\text{poly}(|G|)$, see Wigderson and Xiao [17].)

Throughout, we apply the technique of “derandomized squaring”—analogous to the principal construction in Rozenman and Vadhan’s alternate proof of Reingold’s theorem [15] that Undirected Reachability is in LOGSPACE. In particular, we observe that derandomized squaring provides a generic amplification tool in our setting; specifically, given a constant-bias set S , we can obtain an ε -biased set of size $O(|S|\varepsilon^{-11})$. We also use a tensor product version of derandomized squaring to build ε -biased sets from G recursively, from ε -biased sets for its subgroups or quotients.

Homogeneous direct products and branching programs Groups of the form G^n where G is fixed have been actively studied by the pseudorandomness community as a specialization of the class of constant-width branching programs. The problem of fooling “read-once” group programs induces an alternate notion of ε -biased sets over groups of the form G^n defined by Meka and Zuckerman [10]. Specifically, a read-once branching program on G consists of a tuple $\mathbf{g} = (g_1, \dots, g_n) \in G^n$ and takes a vector of n Boolean variables $\mathbf{b} = (b_1, \dots, b_n)$ as input. At each step, it applies $g_i^{b_i}$, i.e., g_i if $b_i = 1$ and 1 if $b_i = 0$. They say a set $S \subset G^n$ is ε -biased if, for all $\mathbf{b} \neq 0$, the distribution of $\mathbf{g}^{\mathbf{b}}$ is close to uniform, i.e.,

$$\forall h \in G : \left| \Pr_{\mathbf{g} \in S} [\mathbf{g}^{\mathbf{b}} = h] - \frac{1}{|G|} \right| \leq \varepsilon \quad \text{where} \quad \mathbf{g}^{\mathbf{b}} = \prod_{i=1}^n g_i^{b_i}. \quad (1)$$

As they comment, there is no obvious relationship between this definition and the one we consider.¹ We are unable to establish such a connection in general. However, we show in Section 2 that a particular set shown to have property (1) in [10] is also ε -biased in our sense; the proof is completely different. This yields ε -biased sets of size $O(n \cdot \text{poly}(\varepsilon^{-1}))$.

Inhomogeneous direct products For the more general case of groups of the form $G = G_1 \times \dots \times G_n$, we show that a tensor product adaptation of derandomized squaring yields a recursive construction of ε -biased sets of size $\text{poly}(\max_i |G_i|, n, 1/\varepsilon)$.

Normal extensions and “smoothly solvable” groups Finally, we show that if G is solvable and has abelian quotients of bounded exponent, we can construct ε -biased sets of size $(\log |G|)^{1+o(1)} \text{poly}(\varepsilon^{-1})$. Here we use the representation theory of solvable groups to build an ε -biased set for G recursively from those for a normal subgroup H and the quotient G/H .

¹In particular, there is no obvious way to amplify in their setting: for instance, squaring a set S by multiplication in G^n squares the operator norm of any representation, but it has a very complicated effect on the distribution of $\mathbf{g}^{\mathbf{b}}$.

2 An explicit set for G^n with constant ε

Meka and Zuckerman [10] considered the following construction for fooling read-once group branching programs:

Definition 1. Let G be a group and $n \in \mathbb{N}$. Then, given an ε -biased set S over $\mathbb{Z}_{|G|}^n$, define

$$T_S \triangleq \{(g^{s_1}, \dots, g^{s_n}) \mid g \in G, (s_1, \dots, s_n) \in S\}.$$

We prove the following theorem, showing that this construction yields sets of small bias in our sense (and, hence, expander Cayley graphs over G^n).

Theorem 1. *If S is ε -biased over $\mathbb{Z}_{|G|}^n$ then T_S is $(1 - \Omega(1/\log \log |G|)^2 + \varepsilon)$ -biased over G^n .*

Anticipating the proof, we set down the following definition.

Definition 2. Let G be a finite group. For a representation $\rho \in \widehat{G}$ and a subgroup H , define

$$\Pi_H^\rho \triangleq \mathbb{E}_{h \in H} \rho(h)$$

to be the projection operator induced by the subgroup H in ρ . In the case where $H = \langle g \rangle$ is the cyclic group generated by g , we use the following shorthand:

$$\Pi_g^\rho = \Pi_{\langle g \rangle}^\rho.$$

Finally, for groups of the form G^n we use the following convention. Recall that any irreducible representation $\bar{\rho} \in \widehat{G^n}$ is a tensor product, $\bar{\rho} = \bigotimes_{i=1}^n \rho_i$ where $\rho_i \in \widehat{G}$ for each i . That is, if $\bar{g} = (g_1, \dots, g_n)$, then $\bar{\rho}(\bar{g}) = \bigotimes_{i=1}^n \rho_i(g_i)$. Then for an element $g \in G$, we write

$$\Pi_g^{\bar{\rho}} \triangleq \Pi_{\langle g \rangle^n}^{\bar{\rho}} = \bigotimes_{i=1}^n \Pi_g^{\rho_i} \tag{2}$$

for the projection operator determined by the abelian subgroup $\langle g \rangle^n$.

Lemma 2. *Let G be a finite group and ρ a nontrivial irreducible representation of G . Then*

$$\left\| \mathbb{E}_{g \in G} \Pi_g^\rho \right\| \leq 1 - \frac{\phi(|G|)}{|G|} \leq 1 - \Omega\left(\frac{1}{\log \log |G|}\right),$$

where $\phi(\cdot)$ denotes the Euler totient function.

Proof. Expanding the definition of $\Pi_{\langle g \rangle}^\rho$, we have

$$\left\| \mathbb{E}_{g \in G} \Pi_g^\rho \right\| = \left\| \mathbb{E}_{g \in G} \mathbb{E}_{t \in \mathbb{Z}_{|G|}} \rho(g^t) \right\| \leq \mathbb{E}_{t \in \mathbb{Z}_{|G|}} \left\| \mathbb{E}_g \rho(g^t) \right\|.$$

Recall that the function $x \mapsto x^k$ is a bijection in any group G for which $\gcd(|G|, k) = 1$. Moreover, for such k , $\mathbb{E}_g \rho(g^k) = \mathbb{E}_g \rho(g) = 0$ as $\rho \neq 1$. Assuming pessimistically that $\|\mathbb{E}_g \rho(g^k)\| = 1$ for all other k yields the bound $\|\mathbb{E}_{g \in G} \Pi_g^\rho\| \leq 1 - \phi(|G|)/|G|$ promised in the statement of the lemma. The function $\phi(n)$ has the property that

$$\phi(n) > \frac{n}{e^\gamma \log \log n + \frac{3}{\log \log n}}$$

for $n > 3$, where $\gamma \approx .5772\dots$ is the Euler constant [14]; this yields the second estimate in the statement of the lemma. \square

Our proof will rely on the following tail bound for products of operator-valued random variables, proved in Appendix B.

Theorem 3. *Let $\mathbf{P}(H)$ denote the cone of positive operators on the Hilbert space H . Let P_1, \dots, P_k be independent random variables taking values in $\mathbf{P}(H)$ for which $\|P_i\| \leq 1$ and $\|\mathbb{E}[P_i]\| \leq 1 - \delta$. Then*

$$\Pr \left[\|P_k \cdots P_1\| \geq \sqrt{\dim H} \exp \left(-\frac{k\delta}{6} \right) \right] \leq \dim H \cdot \exp \left(-\frac{k\delta^2}{13} \right).$$

We return to the proof of Theorem 1.

Proof of Theorem 1. For a non-trivial irrep $\bar{\rho} = \rho_1 \otimes \cdots \otimes \rho_n \in \widehat{G}^n$, we write

$$\mathbb{E}_{\bar{t} \in T_S} \bar{\rho}(\bar{t}) = \mathbb{E}_{g \in G} \mathbb{E}_{\bar{s} \in S} \bar{\rho}(g^{\bar{s}}) = \mathbb{E}_{g \in G} \mathbb{E}_{\bar{s} \in S} (\text{Res}_{\langle g \rangle^n} \bar{\rho})(g^{\bar{s}}),$$

where $\bar{s} = (s_1, \dots, s_n)$, $g^{\bar{s}} = (g^{s_1}, \dots, g^{s_n})$, and $\text{Res}_H \bar{\rho}$ denotes the restriction of $\bar{\rho}$ to the subgroup $H \subseteq G^n$. For a particular $g \in G$, we decompose the restricted representation $\text{Res}_{\langle g \rangle^n} \bar{\rho}$ into a direct sum of irreps of the abelian group $\langle g \rangle^n \cong \mathbb{Z}_{|\langle g \rangle|}^n$. This yields

$$\text{Res}_{\langle g \rangle^n} \bar{\rho} = \bigoplus_{\chi \in \widehat{\langle g \rangle^n}} \chi^{\oplus a_\chi},$$

where each χ is a one-dimensional representation of the cyclic group $\langle g \rangle^n$ and a_χ denotes the multiplicity with which χ appears in the decomposition.

Now, as S is an ε -biased set over $\mathbb{Z}_{|G|}^n$, its quotient modulo any divisor d of $|G|$ is ε -biased over \mathbb{Z}_d^n . It follows that

$$\left| \mathbb{E}_{\bar{s} \in S} \chi(\bar{s}) \right| \leq \varepsilon$$

for any nontrivial χ ; when χ is trivial, the expectation is 1. Thus for any fixed $g \in G$ we may write

$$\mathbb{E}_{\bar{s} \in S} (\text{Res}_{\langle g \rangle^n} \bar{\rho})(g^{\bar{s}}) = \Pi_g^{\bar{\rho}} + E_g^{\bar{\rho}}.$$

Recall that $\Pi_g^{\bar{\rho}}$ is the projection operator onto the space associated with the copies of the trivial representation of $\langle g \rangle^n$ in $\text{Res}_{\langle g \rangle^n} \bar{\rho}$, i.e., the expectation we would obtain if \bar{s} ranged over all of $\langle g \rangle^n$ instead over just S . The “error operator” $E_g^{\bar{\rho}}$ arises from the nontrivial representations of $\langle g \rangle^n$ appearing in $\text{Res}_{\langle g \rangle^n} \bar{\rho}$, and has operator norm bounded by ε . It follows that

$$\begin{aligned} \left\| \mathbb{E}_{\bar{t} \in T} \bar{\rho}(\bar{t}) \right\| &= \left\| \mathbb{E}_{g \in G} \left(\mathbb{E}_{\bar{s} \in S} \bar{\rho}(g^{\bar{s}}) \right) \right\| = \left\| \mathbb{E}_{g \in G} (\Pi_g^{\bar{\rho}} + E_g^{\bar{\rho}}) \right\| \\ &\leq \left\| \mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right\| + \left\| \mathbb{E}_{g \in G} E_g^{\bar{\rho}} \right\| \leq \left\| \mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right\| + \varepsilon, \end{aligned}$$

and it remains to bound $\left\| \mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right\|$.

As $\mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}}$ is Hermitian, for any positive k we have

$$\left\| \mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right\| = \sqrt[k]{\left\| \left(\mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right)^k \right\|}, \quad (3)$$

so we focus on the operator $\left(\mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right)^k$. Expanding $\Pi_g^{\bar{\rho}} = \bigotimes_i \Pi_g^{\rho_i}$, we may write

$$\left(\mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right)^k = \mathbb{E}_{g_1, \dots, g_k} [\Pi_{g_1}^{\bar{\rho}} \cdots \Pi_{g_k}^{\bar{\rho}}] = \mathbb{E}_{g_1, \dots, g_k} \left[\bigotimes_{i=1}^n \Pi_{g_1}^{\rho_i} \cdots \Pi_{g_k}^{\rho_i} \right]. \quad (4)$$

As $\bar{\rho}$ is nontrivial, there is some coordinate j for which ρ_j is nontrivial. Combining (4) with the fact that $\|A \otimes B\| = \|A\| \|B\|$, we conclude that

$$\left\| \left(\mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right)^k \right\| \leq \mathbb{E}_{g_1, \dots, g_k} \left\| \bigotimes_{i=1}^n \Pi_{g_1}^{\rho_i} \cdots \Pi_{g_k}^{\rho_i} \right\| \leq \mathbb{E}_{g_1, \dots, g_k} \left\| \Pi_{g_1}^{\rho_j} \cdots \Pi_{g_k}^{\rho_j} \right\|. \quad (5)$$

Lemma 2 asserts that $\|\mathbb{E}_g \Pi_g^{\rho_j}\| \leq 1 - \delta_G$, where $\delta_G = \Omega(1/\log \log |G|)$. It follows then from Theorem 3 that

$$\Pr_{g_1, \dots, g_k} \left[\underbrace{\left\| \Pi_{g_1}^{\rho_j} \cdots \Pi_{g_k}^{\rho_j} \right\|}_{(\ddagger)} \geq \sqrt{d_j} \exp(-k\delta_G/6) \right] \leq d_j \cdot \exp(-k\delta_G^2/13), \quad (6)$$

where $d_j = \dim \rho_j$. This immediately provides a bound on $\|(\mathbb{E}_g \Pi_g^{\bar{\rho}})^k\|$. Specifically, combining (5) with (6), let us pessimistically assume that $\|\Pi_{g_1}^{\rho_j} \cdots \Pi_{g_k}^{\rho_j}\| = d_j \exp(-k\delta_G/6)$ for tuples (g_1, \dots, g_k) that do not enjoy property (\ddagger) , and 1 for tuples that do. Then

$$\begin{aligned} \left\| \left(\mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right)^k \right\| &\leq \mathbb{E}_{g_1, \dots, g_k} \left\| \Pi_{g_1}^{\rho_j} \cdots \Pi_{g_k}^{\rho_j} \right\| \\ &\leq d_j \exp(-k\delta_G^2/13) + (1 - d_j \exp(-k\delta_G^2/13)) \sqrt{d_j} \exp(-k\delta_G/6) \\ &\leq 2d_j \exp(-k\delta_G^2/13), \end{aligned}$$

and hence

$$\left\| \mathbb{E}_{g \in G} \Pi_g^{\bar{\rho}} \right\| \leq \inf_k \left(\sqrt[k]{2d_j} \right) \cdot \exp(-\delta_G^2/13) = 1 - \Omega(1/\log \log |G|)^2,$$

where we take the limit of large k . □

3 Derandomized squaring and amplification

In this section we discuss how to amplify ε -biased sets in a generic way. Specifically, we use derandomized squaring to prove the following.

Theorem 4. *Let G be a group and S an $1/10$ -biased set on G . Then for any $\varepsilon > 0$, there is an ε -biased set S_ε on G of size $O(|S|\varepsilon^{-11})$. Moreover, assuming that multiplication can be efficiently implemented in G , the set S_ε can be constructed from S in time polynomial in $|S_\varepsilon|$.*

We have made no attempt to improve the exponent of ε in $|S_\varepsilon|$.

Our approach is similar to [15]. Roughly, if S is an ε -biased set on G we can place a degree- d expander graph Γ on the elements of S to induce a new set

$$S \times_\Gamma S \triangleq \{st \mid (s, t) \text{ an edge of } \Gamma\}.$$

If $\rho : G \rightarrow \mathbf{U}(V)$ is a nontrivial representation of G , by assumption $\|\mathbb{E}_{s \in S} \rho(g)\| \leq \varepsilon$. Applying a natural operator-valued Rayleigh quotient for expander graphs (see Lemma 5 below), we conclude that

$$\left\| \mathbb{E}_{(s,t) \in \Gamma} \rho(s)\rho(t) \right\| = \left\| \mathbb{E}_{(s,t) \in \Gamma} \rho(st) \right\| \leq \lambda(\Gamma) + \varepsilon^2.$$

If Γ comes from a family of Ramanujan-like expanders, then $\lambda(\Gamma) = \Theta(1/\sqrt{d})$, and we can guarantee that $\lambda(\Gamma) = O(\varepsilon^2)$ by selecting $d = \Theta(\varepsilon^{-4})$. The size of the set then grows by a factor of $|S \times_\Gamma S|/|S| = d = \Theta(\varepsilon^{-4})$. We make this precise in Lemma 6 below, which regrettably loses an additional factor of ε^{-1} .

Preparing for the proof of Theorem 4, we record some related material on expander graphs.

Expanders and derandomized products For a d -regular graph $G = (V, E)$, let A denote its normalized adjacency matrix: $A_{uv} = 1/d$ if $(u, v) \in E$ and 0 otherwise. Then A is stochastic, normal, and has operator norm $\|A\| = 1$; the uniform eigenvector \mathbf{y}^+ given by $y_s^+ = 1$ for all $s \in V$ has eigenvalue 1. When G is connected, the eigenspace associated with 1 is spanned by this eigenvector, and all other eigenvalues lie in $[-1, 1)$.

Bipartite graphs will play a special role in our analysis. We write a bipartite graph G on the bipartition U, V as the tuple $G = (U, V; E)$. In a regular bipartite graph, we have $|U| = |V|$ and -1 is an eigenvalue of A associated with the eigenvector \mathbf{y}^- which is $+1$ for $s \in U$ and -1 for $s \in V$. When G is connected, the eigenspace associated with -1 is one-dimensional, and all other eigenvalues lie in $(-1, 1)$: we let $\lambda(G) < 1$ be the leading nontrivial eigenvalue:

$$\lambda(G) = \sup_{\mathbf{y} \perp \mathbf{y}^\pm} \|M\mathbf{y}\| / \|\mathbf{y}\|.$$

When $\mathbf{y} \perp \mathbf{y}^\pm$, observe that $|\langle \mathbf{y}, M\mathbf{y} \rangle| \leq \|\mathbf{y}\| \cdot \|M\mathbf{y}\| \leq \lambda \|\mathbf{y}\|^2$ by Cauchy-Schwarz.

We say that a d -regular, connected, bipartite graph $G = (U, V; E)$ for which $|U| = |V| = n$ and $\lambda(G) \leq \Lambda$ is a *bipartite (n, d, Λ) -expander*. A well-known consequence of expansion is that the ‘‘Rayleigh quotient’’ determined by the expander is bounded: for any function $f : U \cup V \rightarrow \mathbb{R}$ defined on the vertices of a (n, d, λ) expander for which $\sum_{u \in U} f(u) = \sum_{v \in V} f(v) = 0$,

$$\mathbb{E}_{(u,v) \in E} f(u)f(v) \leq \lambda \|f\|_2^2.$$

We will apply a version of this property pertaining to operator-valued functions.

Lemma 5. *Let $G = (U, V; E)$ be a bipartite (n, d, λ) -expander. Associate with each vertex $s \in U \cup V$ a linear operator X_s on the vector space \mathbb{C}^d such that $\|X_s\| \leq 1$, $\|\mathbb{E}_{u \in U} X_u\| \leq \varepsilon_U$, and $\|\mathbb{E}_{v \in V} X_v\| \leq \varepsilon_V$. Then*

$$\left\| \mathbb{E}_{(u,v) \in E} X_u X_v \right\| \leq \lambda + (1 - \lambda) \varepsilon_U \varepsilon_V.$$

We will sometimes apply Lemma 5 to the tensor product of operators. That is, given the same assumptions, we have

$$\left\| \mathbb{E}_{(u,v) \in E} X_u \otimes X_v \right\| \leq \lambda + (1 - \lambda) \varepsilon_U \varepsilon_V.$$

To see this, simply apply the lemma to the operators $X_u \otimes \mathbb{1}$ and $\mathbb{1} \otimes X_v$.

Critical in our setting is the fact that this conclusion is independent of the dimension d . A proof of this folklore lemma appears in Appendix A; see also [6] for a related application to branching programs over groups.

Amplification We return now to the problem of amplifying ε -biased sets over general groups.

Lemma 6. *Let S be an ε -biased set on the group G . Then there is an ε' -biased set S' on G for which $\varepsilon' \leq 5\varepsilon^2$ and $|S'| \leq C|S|\varepsilon^{-5}$, where C is a universal constant. Moreover, assuming that multiplication can be efficiently implemented in G , the set S' can be constructed from S in time polynomial in $|S'|$.*

Proof. We proceed as suggested above. The only wrinkle is that we need to introduce an expander graph on the elements of S that achieves second eigenvalue $\Theta(\varepsilon^2)$.

We apply the explicit family of Ramanujan graphs due to Lubotzky, Phillips, and Sarnak [9]. For each pair of primes p and q congruent to 1 modulo 4, they obtain a graph $\Gamma_{p,q}$ with $p(p^2 - 1)$ vertices, degree $q + 1$, and $\lambda(\Gamma_{p,q}) = 2\sqrt{q}/(q + 1) < 2/\sqrt{q}$. We treat $\Gamma_{p,q}$ as a bipartite graph by taking the double cover: this introduces a pair of vertices, v_A and v_B , for each vertex v of $\Gamma_{p,q}$

and introduces an edge (u_A, v_B) for each edge (u, v) . This graph has eigenvalues $\pm\lambda$ for each eigenvalue λ of $\Gamma_{p,q}$, so except for the ± 1 eigenspace the spectral radius is unchanged.

As we do not have precise control over the number of vertices in this expander family, we will use a larger graph and approximately tile each side with copies of S . Specifically, we select the smallest primes $p, q \equiv 1 \pmod{4}$ for which

$$p(p^2 - 1) > |S| \cdot \lceil \varepsilon^{-1} \rceil \quad \text{and} \quad 2/\sqrt{q} \leq \varepsilon^2. \quad (7)$$

We now associate elements of S with the vertices (of each side) of $\Gamma = \Gamma_{p,q} = (U, V; E)$ as uniformly as possible; specifically, we partition the vertices of U and V into a family of blocks, each of size $|S|$; this leaves a set of less than $|S|$ elements uncovered on each side. Then elements in the blocks are directly associated with elements of S ; the “uncovered” elements may in fact be assigned arbitrarily. As $|U| = |V| \geq |S| \lceil \varepsilon^{-1} \rceil$, the uncovered elements above comprise less than an ε -fraction of the vertices. As above, we define the set $S \times_\Gamma S \triangleq \{uv \mid (u, v) \in E\}$ (where we blur the distinction between a vertex and the element of S to which it has been associated).

Consider, finally, a nontrivial representation ρ of G . As the average over any block of U or V has operator norm no more than ε , and we have an ε -fraction of uncovered elements, the average of ρ over each of U and V is no more than $(1 - \varepsilon)\varepsilon + \varepsilon \leq 2\varepsilon$. Applying Lemma 5, we conclude that $\|\mathbb{E}_{s \in S \times_\Gamma S} \rho(s)\| \leq (2\varepsilon)^2 + \lambda(\Gamma) \leq 5\varepsilon^2$ by our choice of q (the degree less one).

By Dirichlet’s theorem on the density of primes in arithmetic progressions, p and q need be no more than (say) a constant factor larger than the lower bounds $p(p^2 - 1) > |S| \varepsilon^{-1}$ and $q \geq 4\varepsilon^{-4}$ implied by (7). Thus there is a constant C such that $|S'| = p(p^2 - 1)(q + 1) \leq C|S| \cdot \varepsilon^{-5}$. \square

Remarks The construction above is saddled with the tasks of identifying appropriate primes p and q , and constructing the generators for the associated expander of [9]. While these can clearly be carried out in time polynomial in $|S'|$, alternate explicit constructions of expander graphs [12] can significantly reduce this overhead. However, no known explicit family of Ramanujan graphs appears to provide enough density to avoid the tiling construction above. On the other hand, expander graphs with significantly weaker properties would suffice for the construction: any uniform bound of the form $\lambda \leq c\sqrt{\text{degree}}$ would be enough.

Proof of Theorem 4. We apply Lemma 6 iteratively. Set $\varepsilon_0 = 1/10$. After t applications, we have an ε_t -biased set where $\varepsilon_t = 2^{-2^t}/5$. After $t = \lceil \log_2 \log_2(1/5\varepsilon) \rceil$ steps, we have $5\varepsilon^2 \leq \varepsilon_t \leq \varepsilon$. The total increase in size is

$$\frac{|S_\varepsilon|}{|S|} = C^t \left(\prod_{i=0}^{t-1} \varepsilon_i \right)^{-5} = C^t \left(\frac{2\varepsilon_t}{5^{t-1}} \right)^{-5} \leq (C/5)^t (50\varepsilon^2)^{-5} = O(\varepsilon^{-10} (\log \varepsilon^{-1})^{O(1)}) = O(\varepsilon^{-11}).$$

\square

Combining Theorem 4 with the ε -biased sets constructed in Section 2 we establish a family of ε -biased set over G^n for smaller ε :

Theorem 7. *Fix a group G . There is an ε -biased set in G^n of size $O(n\varepsilon^{-11})$ that can be constructed in time polynomial in n and ε^{-1} .*

Proof. Alon et al. [1] construct a families of explicit codes over finite fields which, in particular, offer δ -biased sets over \mathbb{Z}_p^n of size $O(n)$ for any constant δ . As G is fixed, applying Theorem 1 to these sets over $\mathbb{Z}_{|G|}$ with sufficiently small $\delta \approx 1/\log \log |G|$ yields an ε_0 -biased set S_0 over G^n , where ε_0 is a constant close to one (depending on the size of G and the constant δ). We cannot directly apply Theorem 4 to S_0 , as the bias may exceed $1/10$. To bridge this constant gap (from ε_0 to $1/10$), we apply the construction of the proof of Theorem 4 with a slight adaptation. Selecting a small constant α , we may enlarge the expander graph to ensure that it has size at least $|S_0|(1/\alpha)$; then the resulting error guarantee on each side of the graph bipartition is no

more than $\alpha + (1 - \alpha)\varepsilon$ and the product set has bias no more than $(\alpha + \varepsilon)^2 + \lambda(\Gamma)$. This can be brought as close as desired to ε^2 with appropriate selection of the constants α and $\lambda(G)$. As $\lambda(G)$ is constant, this transformation likewise increases the size of the set by a constant, and this method can reduce the error to $1/10$, say, with a constant-factor penalty in the size of S_0 . At this point, Theorem 4 applies, and establishes the bound of the theorem. \square

4 Inhomogeneous direct products

Groups of the form $G = G_1 \times \cdots \times G_n$ appear to frustrate natural attempts to borrow ε -biased sets directly from abelian groups as we did for G^n in Section 2. In this section, we build an ε -biased set for groups of this form by iterating a construction that takes ε -biased sets on two groups G_1 and G_2 and stitches them together, again with an expander graph, to produce an ε' -biased set on $G_1 \times G_2$. In essence, we again use derandomized squaring, but now for the tensor product of two operators rather than their matrix product.

Construction 1. Let G_1 and G_2 be two groups; for each $i = 1, 2$, let S_i be an ε_i -biased set on G_i . We assume that $|S_1| \leq |S_2|$. Let $\Gamma = (U, V; E)$ be a bipartite $(|S_2|, d, \lambda)$ -expander. Associate elements of V with elements of S_2 and, as in the proof of Lemma 6, associate elements of S_1 with U as uniformly as possible. As above, we order the elements of U and tile them with copies of S_1 , leaving a collection of no more than $|S_1|$ vertices “uncovered”; these vertices are then assigned to an initial subset of S_1 of appropriate size. Define $S_1 \otimes_\Gamma S_2 \subset G_1 \times G_2$ to be the set of edges of Γ (realized as group elements according to the association above).

Recall that an irreducible representation ρ of $G_1 \times G_2$ is a tensor product $\rho_1 \otimes \rho_2$, where each ρ_i is an irrep of G_i and $\rho(g_1, g_2) = \rho_1(g_1) \otimes \rho_2(g_2)$. If ρ is nontrivial, then one or both of ρ_1 and ρ_2 is nontrivial, and the bias we achieve on ρ will depend on which of these is the case.

Claim 8. *Assuming that $|S_1| \leq |S_2|$, the set $S_1 \otimes_\Gamma S_2$ of Construction 1 has size $d|S_2|$ and bias no more than*

$$\max \left(\varepsilon_2, \varepsilon_1 + \frac{|S_1|}{|S_2|}, \lambda + \varepsilon_2 \left(\varepsilon_1 + \frac{|S_1|}{|S_2|} \right) \right).$$

Proof. The size bound is immediate. As for the bias, let $\rho = \rho_1 \otimes \rho_2$ be nontrivial. If $\rho_1 = \mathbb{1}$,

$$\left\| \mathbb{E}_{s \in S_1 \otimes_\Gamma S_2} (\rho_1 \otimes \rho_2)(s) \right\| = \left\| \mathbb{E}_{v \in V} \rho_2(v) \right\| \leq \varepsilon_2, \quad (8)$$

as S_2 is in one-to-one correspondence with V . In contrast, if $\rho_2 = \mathbb{1}$, the best we can say is that

$$\left\| \mathbb{E}_{s \in S_1 \otimes_\Gamma S_2} (\rho_1 \otimes \rho_2)(s) \right\| = \left\| \mathbb{E}_{u \in U} \rho_1(u) \right\| \leq \left(1 - \frac{|S_1|}{|S_2|} \right) \varepsilon_1 + \frac{|S_1|}{|S_2|} \leq \varepsilon_1 + \frac{|S_1|}{|S_2|} \quad (9)$$

as in the proof of Lemma 6. When both ρ_i are nontrivial, applying Lemma 5 to (8) and (9) implies that

$$\left\| \mathbb{E}_{s \in S_1 \otimes_\Gamma S_2} (\rho_1 \otimes \rho_2)(s) \right\| \leq \lambda + \varepsilon_2 \left(\varepsilon_1 + \frac{|S_1|}{|S_2|} \right), \quad (10)$$

as desired. \square

Finally, we apply Construction 1 to groups of the form $G_1 \times \cdots \times G_n$.

Theorem 9. *Let $G = G_1 \times \cdots \times G_n$. Then, for any ε , there is an ε -biased set in G of size $\text{poly}(\max_i |G_i|, n, \varepsilon^{-1})$. Furthermore, the set can be constructed in time polynomial in its size.*

Proof. Given the amplification results of Section 3, we may focus on constructing sets of constant bias. We start by adopting the entire group G_i as a 0-biased set for each G_i , and then recursively apply Construction 1. This process will only involve expander graphs of constant degree, which simplifies the task of finding the expander required for Construction 1. In this case, one can construct a constant degree expander graph of desired constant spectral gap on a set X by covering the vertices of X with a family of overlapping expander graphs, uniformizing the degree arbitrarily, and forming a small power of the result. So long as the pairwise intersections of the covering expanders are not too small, the resulting spectral gap can be controlled uniformly. (This luxury was not available to us in the proof of Lemma 6, since in that setting we required λ tending to zero, and insisted on a Ramanujan-like relationship between λ and the degree.)

The recursive construction proceeds by dividing G into two factors: $A = G_1 \times \cdots \times G_{n'}$ and $B = G_{n'+1} \times \cdots \times G_n$, where $n' = \lceil n/2 \rceil$. Given small-biased sets S_A and S_B , we combine them using Construction 1. Examining Claim 8, we wish to ensure that $|S_A|/|S_B|$ is a small enough constant. To arrange for this, we assume without loss of generality that $|S_B| \geq |S_A|$ and duplicate S_B five times, resulting in a (multi-)set S'_B such that $|S_A|/|S'_B| \leq 1/5$.

Assume that each of the recursively constructed sets S_A, S_B has bias at most $1/4$. We apply Construction 1 to S_A and S'_B with an expander Γ of degree d for which $\lambda \leq 1/8$, producing the set $S = S_A \otimes_\Gamma S'_B$. Ideally, we would like S to also be $1/4$ -biased, in which case a set of constant bias and size $\text{poly}(\max_i |G_i|, n)$ would follow by induction.

Let $\rho = \rho_A \otimes \rho_B$ be nontrivial, where $\rho_A \in \hat{A}$ and $\rho_B \in \hat{B}$. If $\rho_A = \mathbb{1}$ then, as in (8), $\|\mathbb{E}_{s \in S} \rho(s)\| \leq 1/4$. Likewise, if both ρ_A and ρ_B are nontrivial, (10) gives $\|\mathbb{E}_{s \in S} \rho(s)\| \leq 1/8 + 1/4(1/4 + 1/5) \leq 1/4$. At first inspection, the case where $\rho_B = \mathbb{1}$ appears problematic, as (9) only provides the discouraging estimate $\|\mathbb{E}_{s \in S} \rho(s)\| \leq 1/4 + 1/5$. Thus it seems possible that iterative application of Construction 1 could lose control of the error. However, as long as the tiling of U , the left side of the expander in Construction 1, is carried out in a way that ensures that the uncovered elements of U are tiled with respect to *previous* stages of the recursive construction, it is easy to check that subsequent recursive appearances of this case can contribute no more than the geometric series $1/5 + (1/5)^2 + \cdots = 1/4$ to the bias. Any following recursive application of the construction in which the representation is nontrivial in both blocks will then drive the error back to $1/4$, as $1/8 + 1/4(1/4 + 1/4) = 1/4$. (If this case occurs at the last stage of recursion, then S still has bias at most $1/4 + 1/5 \leq 1/2$.)

Recall that for the base case of the induction, we treat each G_i as a 0-biased set for itself. Since there are $\log_2 n$ layers of recursion, and each layer multiplies the size of the set by the constant factor $5d$, we end with a $1/2$ -biased set S of size at most $(5d)^{\log_2 n} \max_i |G_i| = \text{poly}(\max_i |G_i|, n)$. Finally, applying the amplification of Theorem 4, after first driving the bias down to $1/10$ as in Theorem 7, completes the proof. \square

We note that if the G_i are of polynomial size, then we can use the results of Wigderson and Xiao [17] to find ε -biased sets of size $O(\log |G_i|)$ in time $\text{poly}(|G_i|)$. Using these sets in the base case of our recursion then gives a ε -biased set for G of size $\text{poly}(\max_i \log |G_i|, n, \varepsilon^{-1})$.

5 Normal extensions and smoothly solvable groups

While applying these techniques to arbitrary groups (even in the case when they have plentiful subgroups) seems difficult, for solvable groups can again use a form of derandomized squaring. First, recall the derived series: if G is solvable, then setting $G^{(0)} = G$ and taking commutator subgroups $G^{(i+1)} = [G^{(i)}, G^{(i)}]$ gives a series of normal subgroups,

$$1 = G^{(\ell)} \triangleleft \cdots \triangleleft G^{(1)} \triangleleft G^{(0)} = G.$$

We say that ℓ is the *derived length* of G . Each factor $G^{(i)}/G^{(i+1)} = A_i$ is abelian, and $G^{(i)}$ is normal in G for all i . Since $|A_i| \geq 2$, it is obvious that $\ell = O(\log |G|)$. However, more is

true. The *composition series* is a refinement of the derived series where each quotient is a cyclic group of prime order, and the length c of this refined series is the *composition length*. Clearly $c \leq \log_2 |G|$. Glasby [8] showed that $\ell \leq 3 \log_2 c + 9 = O(\log c)$, so $\ell = O(\log \log |G|)$.

We focus on groups that are *smoothly solvable* [7], in the sense that the abelian factors have constant exponent. (Their definition of smooth solvability allows the factors to be somewhat more general, but we avoid that here for simplicity.) We then have the following:

Theorem 10. *Let G be a solvable group, and let its abelian factors be of the form $A_i = \mathbb{Z}_{p_i}^t$ (or factors of such groups) where $p_i = O(1)$. Then G possesses an ε -biased set S_ε of size $(\log |G|)^{1+o(1)} \text{poly}(\varepsilon^{-1})$.*

We deliberately gloss over the issue of explicitness. However, we claim that if G is polynomially uniform in the sense of [11], so that we can efficiently express group elements and products as a string of coset representatives in the derived series, then S_ε can be computed in time polynomial in its size.

Proof. Solvable groups can be approached via Clifford theory, which controls the structure of representations of a group G when restricted to a normal subgroup. In fact, we require only a simple fact about this setting. Namely, if $H \triangleleft G$ and ρ is an irrep of G , then either $\text{Res}_H \rho$ contains only copies of the trivial representation so that $\rho(h) = \mathbf{1}_{\rho_d}$ for all $h \in H$, or $\text{Res}_H \rho$ contains *no* copies of the trivial representation.

It is easy to see that the irreps ρ of G for which $\text{Res}_H \rho$ is trivial are in one-to-one correspondence with irreps of the group G/H , and we will blur this distinction. With this perspective, it is natural to attempt to assemble an ε -biased set for G from S_H , an ε_H -biased set for H , and $S_{G/H}$, an $\varepsilon_{G/H}$ -biased set for G/H . While $S_H \subset H \subset G$, there is—in general—no subgroup of G isomorphic to G/H , so it is not clear how to appropriately embed $S_{G/H}$ into G . Happily, we will see that reasonable bounds can be obtained even with an arbitrary embedding. In particular, we treat $S_{G/H}$ as a subset of G by lifting each element $x \in S_{G/H}$ to an arbitrary element $\hat{x} \in G$ lying in the H -coset associated with x .

If S_H and $S_{G/H}$ were the same size, and we could directly introduce an expander graph Γ on $S_H \times S_{G/H}$, then Lemma 5 could still be used to control the bias of $S = \{s\hat{t} \mid (s, t) \in \Gamma\}$. Specifically, consider a nontrivial representation ρ of G . If $\text{Res}_H \rho$ is trivial, then analogous to (8) we have $\|\mathbb{E}_{s \in S} \rho(s)\| = \|\mathbb{E}_{s \in S_{G/H}} \rho(s)\| \leq \varepsilon_{G/H}$. On the other hand, if $\text{Res}_H \rho$ restricts to H without any appearances of the trivial representation, then $\|\mathbb{E}_{h \in S_H} \rho(h)\| \leq \varepsilon_H$. In this case, the action of the elements of $S_{G/H}$ on ρ may be quite pathological, permuting and “twiddling” the H -irreps appearing in $\text{Res}_H \rho$. However, as $\|\rho(s)\| = 1$ (by unitarity) for all $s \in S_{G/H}$, we can conclude from Lemma 5 that $\|\mathbb{E}_{s \in S} \rho(s)\| \leq \lambda(\Gamma) + \varepsilon_H$.

We recursively apply the construction outlined above, accounting for the “tiling error” of finding an appropriate expander. Specifically, let us inductively assume we have ϵ -biased sets S^+ on $G^+ = G/G^{(k)}$ and S^- on $G^- = G^{(k)}$ for $k = \lceil \ell/2 \rceil$, where ℓ is the derived length of G . Selecting an expander graph Γ of size at least $\alpha^{-1} \max(|S^-|, |S^+|)$ and $\lambda(\Gamma) \leq \alpha$, for an α to be determined, we tile each side of the graph with elements from S^- and S^+ , completing them arbitrarily on the “uncovered elements.” Since at most a fraction α of the elements on either side are uncovered, the average of a nontrivial representation over either side of the expander has operator norm no more than $\epsilon + \alpha$. Lemma 5 then implies that the bias of the set $S = \{s\hat{t} \mid (s, t) \in \Gamma\}$ is at most $\lambda(\Gamma) + (\epsilon + \alpha) \leq \epsilon + 2\alpha$. If we use the Ramanujan graphs of [9] described above, we can achieve degree $O(\alpha^{-2})$ and size $O(\alpha \max(|S^-|, |S^+|))$. Thus, each recursive step of this process scales the sizes of the sets by a factor $O(\alpha^{-3})$ and introduces additive error 2α . The number of levels of recursion is $\lceil \log_2 \ell \rceil$, so if we choose $\alpha < 1/(4 \lceil \log \ell \rceil)$ then the total accumulated error is less than $1/2$.

Assuming that we have α -biased sets for each abelian factor A_i of size no more than s , this yields a $1/2$ -biased set S for G of size $s\alpha^{-3 \log_2 \ell} = s(\log \ell)^{O(\log \ell)}$. For constant p , there

are α -biased sets for \mathbb{Z}_p^n [1] of size $s = O(n/\alpha^3) = (\log |G|)(\log \ell)^{O(1)}$. Using the fact [8] that $\ell = O(\log \log |G|)$, the total size of S is

$$(\log |G|)(\log \ell)^{O(\log \ell)} = (\log |G|)(\log \log \log |G|)^{O(\log \log \log |G|)} = (\log |G|)^{1+o(1)}.$$

Finally, we amplify S to an ε -biased set S_ε for whatever ε we desire with Theorem 4, introducing a factor $O(\varepsilon^{-11})$. \square

Acknowledgments

We thank Amnon Ta-Shma, Emanuele Viola, and Avi Wigderson for helpful discussions. This work was supported by NSF grant CCF-1117426 and ARO contract W911NF-04-R-0009.

References

- [1] N. Alon, J. Bruck, J. Naor, M. Naor, and R. M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Transactions on Information Theory*, 38(2):509–516, 1992. doi:[10.1109/18.119713](https://doi.org/10.1109/18.119713).
- [2] Noga Alon and Yuval Roichman. Random Cayley graphs and expanders. *Random Structures and Algorithms*, 5(2):271–284, 1994. doi:[10.1002/rsa.3240050203](https://doi.org/10.1002/rsa.3240050203).
- [3] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple construction of almost k-wise independent random variables. *Random Struct. Algorithms*, 3(3):289–304, 1992.
- [4] Avraham Ben-Aroya and Amnon Ta-Shma. Constructing small-bias sets from algebraic-geometric codes. *Theory of Computing*, 9(5):253–272, 2013. doi:[10.4086/toc.2013.v009a005](https://doi.org/10.4086/toc.2013.v009a005).
- [5] A. Bogdanov and E. Viola. Pseudorandom bits for polynomials. *SIAM Journal on Computing*, 39(6):2464–2486, 2010. doi:[10.1137/070712109](https://doi.org/10.1137/070712109).
- [6] A. De. Pseudorandomness for permutation and regular branching programs. In *2011 IEEE 26th Annual Conference on Computational Complexity (CCC)*, pages 221–231, 2011. doi:[10.1109/CCC.2011.23](https://doi.org/10.1109/CCC.2011.23).
- [7] Katalin Friedl, Gábor Ivanyos, Frédéric Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 1–9, 2003.
- [8] S. P. Glasby. The composition and derived lengths of a soluble group. *J. Algebra*, 120:406–413, 1989.
- [9] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. doi:[10.1007/BF02126799](https://doi.org/10.1007/BF02126799).
- [10] Raghu Meka and David Zuckerman. Small-bias spaces for group products. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 5687 of *Lecture Notes in Computer Science*, pages 658–672. Springer, 2009. doi:[10.1007/978-3-642-03685-9_49](https://doi.org/10.1007/978-3-642-03685-9_49).
- [11] Cristopher Moore, Daniel N. Rockmore, and Alexander Russell. Generic quantum fourier transforms. *ACM Transactions on Algorithms*, 2(4):707–723, 2006.

- [12] M. Morgenstern. Existence and explicit constructions of $q + 1$ regular Ramanujan graphs for every prime power q . *Journal of Combinatorial Theory, Series B*, 62(1):44–62, 1994. doi:[10.1006/jctb.1994.1054](https://doi.org/10.1006/jctb.1994.1054).
- [13] J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993. doi:[10.1137/0222053](https://doi.org/10.1137/0222053).
- [14] J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois Journal of Mathematics*, 6:64–94, 1962.
- [15] Eyal Rozenman and Salil Vadhan. Derandomized squaring of graphs. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, volume 3624 of *Lecture Notes in Computer Science*, pages 436–447. Springer, 2005. doi:[10.1007/11538462_37](https://doi.org/10.1007/11538462_37).
- [16] Eugene Schenkman. *Group Theory*. R. E. Krieger Pub. Co., 1975.
- [17] Avi Wigderson and David Xiao. Derandomizing the Ahlswede-Winter matrix-valued Chernoff bound using pessimistic estimators, and applications. *Theory of Computing*, 4(3): 53–76, 2008. doi:[10.4086/toc.2008.v004a003](https://doi.org/10.4086/toc.2008.v004a003).

A Quadratic forms associated with expander graphs

Our goal is to establish the two generalized Rayleigh quotient bounds described in Lemmas 12 and 5. We begin with the following preparatory lemma.

Lemma 11. *Let $G = (U, V; E)$ be a (n, d, λ) -expander. Associate with each vertex $s \in U \cup V$ a vector \mathbf{x}^s in \mathbb{C}^d such that $\mathbb{E}_{u \in U} \mathbf{x}^u = 0$ and $\mathbb{E}_{v \in V} \mathbf{x}^v = 0$. Then*

$$\left| \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^u, \mathbf{x}^v \rangle \right| \leq \lambda \mathbb{E}_s \|\mathbf{x}^s\|^2.$$

Proof. Let X denote the $2n \times d$ matrix whose entries are $X_{sk} = x_k^s$. Then the rows of X are the vectors \mathbf{x} ; for an column index $k \in \{1, \dots, d\}$, we let $\mathbf{y}^k \in \mathbb{C}^{2n}$ denote the vector associated with this column:

$$y_v^k = x_k^v.$$

Considering that $\sum_u \mathbf{x}^u = \sum_v \mathbf{x}^v = 0$, each \mathbf{y}^k is orthogonal to both \mathbf{y}^+ and \mathbf{y}^- .

The expectation over a random edge (u, v) of $\langle \mathbf{x}^u, \mathbf{x}^v \rangle$ can be written

$$\begin{aligned} \left| \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^u, \mathbf{x}^v \rangle \right| &= \left| \mathbb{E}_{(u,v) \in E} \sum_k X_{uk} X_{vk} \right| = \left| \sum_k \mathbb{E}_{(u,v) \in E} X_{uk} X_{vk} \right| \\ &= \left| \sum_k \frac{1}{nd} \sum_{(u,v) \in E} x_k^u x_k^v \right| = \left| \frac{1}{n} \sum_k \frac{1}{d} \sum_{(u,v) \in E} y_u^k y_v^k \right| \\ &= \frac{1}{2n} \left| \sum_k \langle \mathbf{y}^k, A \mathbf{y}^k \rangle \right| \leq \frac{1}{2n} \sum_k |\langle \mathbf{y}^k, A \mathbf{y}^k \rangle| \\ &\leq \frac{\lambda}{2n} \sum_k \|\mathbf{y}^k\|^2 = \frac{\lambda}{2n} \sum_s \|\mathbf{x}^s\|^2 = \lambda \mathbb{E}_s \|\mathbf{x}^s\|^2. \quad \square \end{aligned}$$

Lemma 12. Let $G = (U, V; E)$ be a (n, d, λ) -expander. Associate with each vertex $s \in U \cup V$ a vector \mathbf{x}^s in \mathbb{C}^d such that $\|\mathbb{E}_{u \in U} \mathbf{x}^u\| = \varepsilon_U$ and $\|\mathbb{E}_{v \in V} \mathbf{x}^v\| = \varepsilon_V$. Then

$$\left| \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^u, \mathbf{x}^v \rangle \right| \leq \lambda \left(\mathbb{E}_s \|\mathbf{x}^s\|^2 - \frac{\varepsilon_U^2}{2} - \frac{\varepsilon_V^2}{2} \right) + \varepsilon_U \varepsilon_V.$$

Proof of Lemma 12. Let $\mathbf{x}^U = \mathbb{E}_{u \in U} \mathbf{x}^u$ and $\mathbf{x}^V = \mathbb{E}_{v \in V} \mathbf{x}^v$. We have

$$\left| \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^u, \mathbf{x}^v \rangle \right| = \left| \mathbb{E}_{(u,v) \in E} \langle (\mathbf{x}^u - \mathbf{x}^U) + \mathbf{x}^U, (\mathbf{x}^v - \mathbf{x}^V) + \mathbf{x}^V \rangle \right|$$

which we may further expand into

$$\left| \mathbb{E}_{(u,v) \in E} \langle (\mathbf{x}^u - \mathbf{x}^U), (\mathbf{x}^v - \mathbf{x}^V) \rangle + \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^U, (\mathbf{x}^v - \mathbf{x}^V) \rangle + \mathbb{E}_{(u,v) \in E} \langle (\mathbf{x}^u - \mathbf{x}^U), \mathbf{x}^V \rangle + \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^U, \mathbf{x}^V \rangle \right|. \quad (11)$$

As G is regular, the vertices of a uniformly random edge (u, v) are individually uniform on U and V , from which it follows that the two middle terms of (11) are both zero. Hence we conclude that

$$\left| \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^u, \mathbf{x}^v \rangle \right| \leq \left| \mathbb{E}_{(u,v) \in E} \langle (\mathbf{x}^u - \mathbf{x}^U), (\mathbf{x}^v - \mathbf{x}^V) \rangle \right| + |\langle \mathbf{x}^U, \mathbf{x}^V \rangle|.$$

Applying Lemma 11 to the the vectors

$$\mathbf{x}^u - \mathbf{x}^U \quad \text{and} \quad \mathbf{x}^v - \mathbf{x}^V,$$

we conclude that

$$\left| \mathbb{E}_{(u,v) \in E} \langle (\mathbf{x}^u - \mathbf{x}^U), (\mathbf{x}^v - \mathbf{x}^V) \rangle \right| \leq \frac{\lambda}{2n} \left(\sum_u \|\mathbf{x}^u - \mathbf{x}^U\|^2 + \sum_v \|\mathbf{x}^v - \mathbf{x}^V\|^2 \right).$$

The summation $\sum_u \|\mathbf{x}^u - \mathbf{x}^U\|^2$ can be calculated as follows.

$$\begin{aligned} \sum_u \|\mathbf{x}^u - \mathbf{x}^U\|^2 &= \sum_u \langle \mathbf{x}^u - \mathbf{x}^U, \mathbf{x}^u - \mathbf{x}^U \rangle \\ &= \sum_u (\langle \mathbf{x}^u, \mathbf{x}^u \rangle - \langle \mathbf{x}^u, \mathbf{x}^U \rangle - \langle \mathbf{x}^U, \mathbf{x}^u \rangle + \langle \mathbf{x}^U, \mathbf{x}^U \rangle) \\ &= \sum_u \langle \mathbf{x}^u, \mathbf{x}^u \rangle - n \langle \mathbf{x}^U, \mathbf{x}^U \rangle - n \langle \mathbf{x}^U, \mathbf{x}^U \rangle + n \langle \mathbf{x}^U, \mathbf{x}^U \rangle \\ &= \sum_u \|\mathbf{x}^u\|^2 - n \|\mathbf{x}^U\|^2 \\ &= \sum_u \|\mathbf{x}^u\|^2 - n \varepsilon_U^2. \end{aligned}$$

Therefore,

$$\begin{aligned} \left| \mathbb{E}_{(u,v) \in E} \langle (\mathbf{x}^u - \mathbf{x}^U), (\mathbf{x}^v - \mathbf{x}^V) \rangle \right| &\leq \frac{\lambda}{2n} \left(\sum_u \|\mathbf{x}^u\|^2 - n \varepsilon_U^2 + \sum_v \|\mathbf{x}^v\|^2 - n \varepsilon_V^2 \right) \\ &\leq \lambda \left(\mathbb{E}_s \|\mathbf{x}^s\|^2 - \frac{\varepsilon_U^2}{2} - \frac{\varepsilon_V^2}{2} \right). \end{aligned}$$

By Cauchy-Schwarz, we have $|\langle \mathbf{x}^U, \mathbf{x}^V \rangle| \leq \varepsilon_U \varepsilon_V$. In total, then,

$$\left| \mathbb{E}_{(u,v) \in E} \langle \mathbf{x}^u, \mathbf{x}^v \rangle \right| \leq \lambda \left(\mathbb{E}_s \|\mathbf{x}^s\|^2 - \frac{\varepsilon_U^2}{2} - \frac{\varepsilon_V^2}{2} \right) + \varepsilon_U \varepsilon_V,$$

as desired. \square

Lemma (Restatement of Lemma 5). *Let $G = (U \cup V, E)$ be a (n, d, λ) -expander. Associate with each vertex $s \in U \cup V$ a linear operator X_s on the vector space \mathbb{C}^d such that $\|X_s\| \leq 1$, $\|\mathbb{E}_{u \in U} X_u\| = \varepsilon_U$, and $\|\mathbb{E}_{v \in V} X_v\| = \varepsilon_V$. Then*

$$\left\| \mathbb{E}_{(u,v) \in E} X_u X_v \right\| \leq \lambda + (1 - \lambda) \varepsilon_U \varepsilon_V.$$

Proof of Lemma 5. Let X denote the linear operator $\mathbb{E}_{(u,v) \in E} X_u X_v$. Writing

$$\|X\| = \max_{\substack{\|\mathbf{x}\|=1, \\ \|\mathbf{y}\|=1}} |\langle \mathbf{x}, X\mathbf{y} \rangle|,$$

we observe that

$$\langle \mathbf{x}, X\mathbf{y} \rangle = \left\langle \mathbf{x}, \mathbb{E}_{(u,v) \in E} X_u X_v \mathbf{y} \right\rangle = \mathbb{E}_{(u,v) \in E} \langle X_u^\dagger \mathbf{x}, X_v \mathbf{y} \rangle.$$

Considering the bounds on $\mathbb{E}_u X_u$ and $\mathbb{E}_v X_v$, it follows that $\|\mathbb{E}_u X_u^\dagger \mathbf{x}\| \leq \varepsilon_U$ and $\|\mathbb{E}_v X_v \mathbf{y}\| \leq \varepsilon_V$; applying Lemma 12 with the vector family $\mathbf{x}^u = X_u^\dagger \mathbf{x}$ and $\mathbf{x}^v = X_v \mathbf{y}$ we conclude that

$$\begin{aligned} |\langle \mathbf{x}, X\mathbf{y} \rangle| &\leq \max_{\substack{\delta_U \leq \varepsilon_U \\ \delta_V \leq \varepsilon_V}} \lambda \left(\mathbb{E}_s \|\mathbf{x}^s\|^2 - \frac{\delta_U^2}{2} - \frac{\delta_V^2}{2} \right) + \delta_U \delta_V \\ &\leq \max_{\substack{\delta_U \leq \varepsilon_U \\ \delta_V \leq \varepsilon_V}} \lambda (1 - \delta_U \delta_V) + \delta_U \delta_V \leq \lambda + (1 - \lambda) \varepsilon_U \varepsilon_V \end{aligned}$$

as $\delta_U^2 + \delta_V^2 \geq 2\delta_U \delta_V$. □

B A tail bound for products of operator-valued random variables

Our goal is to establish the following tail bound (a restatement and expansion of Theorem 3).

Theorem (Restatement of Theorem 3). *Let $\mathbf{P}(H)$ denote the cone of positive operators on the Hilbert space H and let P_1, \dots, P_k be independent random variables taking values in $\mathbf{P}(H)$ for which*

$$\|P_i\| \leq 1, \quad \text{and} \quad \|\mathbb{E}[P_i]\| \leq 1 - \delta.$$

Then for any $\Delta \geq 0$,

$$\Pr \left[\|P_k \cdots P_1\| \geq \sqrt{\dim H} \exp \left(-\frac{k\delta}{2} + \Delta \right) \right] \leq \dim H \cdot \exp \left(-\frac{\Delta^2}{2k \ln 2} \right).$$

In particular, choosing $\Delta = k\delta/3$, we conclude that

$$\Pr \left[\|P_k \cdots P_1\| \geq \sqrt{\dim H} \exp \left(-\frac{k\delta}{6} \right) \right] \leq \dim H \cdot \exp \left(-\frac{k\delta^2}{18 \ln 2} \right) \leq \dim H \cdot \exp \left(-\frac{k\delta^2}{13} \right).$$

Recall Azuma's inequality for supermartingales:

Theorem 13 (Azuma's inequality). *Let X_0, \dots, X_T be a family of real-valued random variables for which $|X_i - X_{i-1}| \leq \alpha_i$ and $\mathbb{E}[X_i \mid X_1, \dots, X_{i-1}] \leq X_{i-1}$. Then*

$$\Pr[X_T - X_0 \geq \lambda] \leq \exp \left(-\frac{\lambda^2}{2 \sum_i \alpha_i} \right).$$

Corollary 14. Let X_0, \dots, X_T be a family of real-valued random variables for which $X_{i-1} - \alpha_i \leq X_i \leq X_{i-1}$ and $\mathbb{E}[X_i \mid X_1, \dots, X_{i-1}] \leq X_{i-1} - \varepsilon_i$ for some $\varepsilon_i \leq \alpha_i$. Then

$$\Pr[X_T - X_0 \geq -\sum_i \varepsilon_i + \lambda] \leq \exp\left(-\frac{\lambda^2}{2\sum_i \alpha_i}\right).$$

Proof. Apply Azuma's inequality to the random variables $\tilde{X}_t = X_t + \sum_i^t \varepsilon_i$. \square

Proof of Theorem 3. We begin by considering the behavior of the operator $P_k \cdots P_1$ on a particular vector \mathbf{v} . To complete the proof we will select an orthonormal basis \mathcal{B} of H . The operator norm is bounded above by the Frobenius norm,

$$\|P_k \cdots P_1\| \leq \|P_k \cdots P_1\|_F = \sqrt{\sum_{\mathbf{b} \in \mathcal{B}} \|P_k \cdots P_1 \mathbf{b}\|^2} \leq \sqrt{\dim H} \cdot \max_{\mathbf{b} \in \mathcal{B}} \|P_k \cdots P_1 \mathbf{b}\|. \quad (12)$$

Now fix a unit-length vector $\mathbf{v} \in H$ and consider the random variables

$$\mathbf{v}_0 = \mathbf{v}, \quad \mathbf{v}_1 = P_1 \mathbf{v}, \quad \mathbf{v}_2 = P_2 P_1 \mathbf{v},$$

and

$$\ell_i = \begin{cases} \|\mathbf{v}_i\|/\|\mathbf{v}_{i-1}\| & \text{if } \mathbf{v}_{i-1} \neq 0, \\ 0 & \text{otherwise,} \end{cases}$$

Our goal is to establish strong tail bounds on the random variable $\|\mathbf{v}_k\| = \ell_k \ell_{k-1} \cdots \ell_1$. Recalling that $\|\mathbb{E}[P_i]\| \leq 1 - \varepsilon$ and that the P_i are independent we have

$$\mathbb{E}[\ell_i \mid P_1, \dots, P_{i-1}] \leq 1 - \varepsilon, \quad (13)$$

and we proceed to apply a martingale tail bound.

It will be more convenient to work with log-bounded random variables, so we define $m_i = \max(\ell_i, 1/2)$ and observe that $\|\mathbf{v}_k\| \leq m_k m_{k-1} \cdots m_1$ and $\ln \|\mathbf{v}_k\| \leq \sum_i \ln m_i$. Considering that $\max(x, 1/2) \leq (1+x)/2$ for $x \in [0, 1]$ we conclude from equation (13) above that $\mathbb{E}[m_i \mid P_1, \dots, P_{i-1}] \leq 1 - \varepsilon/2$. Since $1/2 \leq m_i \leq 1$ and $\ln m \leq m - 1$ for $m > 0$, we have

$$\mathbb{E}[\ln m_i \mid P_1, \dots, P_{i-1}] \leq -\varepsilon/2. \quad (14)$$

Applying Azuma's inequality (specifically, Corollary 14 above) to the random variables $M_t = \sum_{i=1}^t \ln m_i$, we conclude that

$$\Pr\left[M_k \geq -\frac{k\varepsilon}{2} + \Delta\right] = \Pr\left[\sum_i \ln m_i \geq -\frac{k\varepsilon}{2} + \Delta\right] \leq \exp\left(-\frac{\Delta^2}{2k \ln 2}\right)$$

and hence

$$\Pr\left[\|P_k \cdots P_1 \mathbf{v}\| \geq \exp\left(-\frac{k\varepsilon}{2} + \Delta\right)\right] \leq \exp\left(-\frac{\Delta^2}{2k \ln 2}\right).$$

Applying the above inequality to an orthonormal basis $\mathbf{b}_1, \dots, \mathbf{b}_n$ of H , we find that

$$\Pr\left[\exists i : \|P_k \cdots P_1 \mathbf{b}_i\|_2 \geq \exp\left(-\frac{k\varepsilon}{2} + \Delta\right)\right] \leq \dim H \cdot \exp\left(-\frac{\Delta^2}{2k \ln 2}\right)$$

by the union bound. Applying (12) then gives

$$\Pr\left[\|P_k \cdots P_1\| \geq \sqrt{\dim H} \exp\left(-\frac{k\varepsilon}{2} + \Delta\right)\right] \leq \dim H \cdot \exp\left(-\frac{\Delta^2}{2k \ln 2}\right). \quad \square$$